SNS Low Level RF Control System Team Response to the Reviewer Comments
of the Dec. 17 Hardware Design Review

31 Jan 2003

The LLRF Team has prepared this written response to the Reviewer Comments that were received following the Dec. 17 hardware design review at Los Alamos. The reviewer comments have been grouped in six categories. The persons responsible for addressing the comments are identified under the category headings, with the principal author's name underlined. The text is color coded according to the following color key:

Color Key
Hovater : Red
Doolittle : Green
Ziomek : Blue
LLRF Team : Black

---

1. GENERAL COMMENTS

Champion, Ratti, Shoaee

Attached (Larry's are below) you will find the design review comments from Larry Doolittle, Chris Ziomek and myself. The overall consensus is to proceed with the hardware design while paying close attention to our individual recommendations and suggestions. A theme that ran through all three reviewer's comments (and a recurring discussion of the review board) is the hardware specifications. I think Larry said it best, "There is a project-wide need to finish the flow-down from accelerator specs (amplitude and phase errors) to hardware specifications." We highly recommend that this be completed by the end of January.

The hardware conceptual design is fairly straight forward and on track. It appears more than adequate for implementing all of the necessary control algorithms especially if the team follows the strategy used in the LBL system. At this stage I would say go for it and build the hardware. Any problems can be worked out on the next rev (it's more important to get some hardware and commissioning under the teams belt).

A concern I (and others) have is the system specifications. We don't have an integrated specification especially as it flows from the golden grail requirements of 1% - 1^{o}. The system/board specifications seem to me good educated guesses for some items and loosely tied back to the overall field control specifications. While it is my belief that the system will be capable of meeting its overall specifications this has not been demonstrated.

Agreed. We are revisiting the system error budget at the subsystem and component levels. Our goal is to convince ourselves and others that the system is fundamentally capable of meeting the high level performance requirements. Hengjie Ma has put together a spreadheet which is a good start.

Lab interplay also seems to be still rough. The system is being built across laboratory lines and comments such as "that is their problem not ours" still exist. I was surprised that the cavity down converter mixer was not immediately known or included in this review. The reference and distribution were also excluded and seeing that the jitter and drift plays a huge role in this it can't be ignored nor is it fair to the people designing the hardware not to have a known jitter budget (see previous paragraph). The review was almost exclusively dominated by a "LANL view" with out much input from LBL or ORNL. I thought the intention from the get go was to grow from the LBL experience. It is obvious that the hardware design (block down converter, single FPGA) has been guided by this strategy, but it is not clear if the firmware is proceeding along the same path.

The LLRF Team was reorganized following a difficult review last September. It has not been an easy transition, and we've made numerous corrections along the way. In spite of these difficulties, the team members are committed to putting together a system that will meet the needs of the SNS accelerator.

One thing that was immediately apparent is that there is really no over all system engineer leading the effort. The three laboratory team leaders Mark, Hamid and Alex do not have the experience or expertise to make the decisions as it applies to the hardware or firmware. This in my mind has hindered the speed of the development especially as it applies to certain design choices. One example where a system engineer could have helped is the FIR filter. Questions have repeatedly been raised about this and better ways suggested, but know one has sat down and looked at other options. Another area is the model simulation and firmware, it is not clear to me that these efforts will converge in the FPGA. I don't think you can implement the model into the system with out some how tying these two together. In addition questions during the review with vague or incomplete answers made it obvious that a lead system engineer is desperately needed.

Larry Doolittle is the lead system engineer per agreement between Mark, Hamid, Alex and Larry. Mark Champion, Hengjie Ma and Amy Regan will be heavily involved in the system engineering as well.

Overall I have good confidence the hardware team will deliver a system. The schedule does seem optimistic especially since the team (Power and Stettler) are also preoccupied with diagnostics projects. What you do with this hardware is another matter. I am concerned that with out a lead system engineer who understands signal processing and control systems that the hardware will not be implemented correctly or efficiently. Given that I would start with the LBL firmware (that's what was sold to the review board) and go from there. Finally I recommend a lead system engineer be chosen (assigned) for all hardware and firmware decisions.

General comments:

* Much progress has been made. There's still a lot of work to be done.

* I see too much emphasis on telling other people what to do, and not enough buckling down and doing the work.

* I have some concern about over-reliance on complex black-box software. If you can't judge whether the results of a computation make physical sense, using your brain and the back of an envelope, you haven't solved the problem.

* As I said in the meeting, as a clearly dissenting opinion, the team should standardize on standards, and buy products because they implement the standards well. Standardizing on products is short-sighted, closed-minded, and (apparently) needlessly expensive.

* There is a project-wide need to finish the flow-down from accelerator specs (amplitude and phase errors) to hardware specifications. I would almost ignore the digital control contribution at this point, other than giving it a non-zero budget. Key items are thermal drift (cables, mixers, amplifiers, ADCs), plus mixer distortion.

5. **Digital Processing** The FPGA development is the critical path item and could possibly significantly delay the delivery of the system. Also, I believe that the overall system performance will ultimately be limited by the FIR filter. I suggest the following:
    a. Leverage the work done at BNL as much as possible.
    b. Use a standard set of implementation tools, language, and style guidelines to ensure compatibility and code reuse amongst the many code developers.
    c. Look at alternatives to the FIR design. The 1.6 microsecond group delay in not acceptable. My experience is that a 50 kHz bandwidth is not sufficient. For example, if your klystron ripple is at 20 kHz, you will have very little gain and noise rejection at this frequency. At a minimum, a non-symmetric, non-linear phase FIR filter will cut the delay in half. Having a mode filter in the feedback path does not eliminate the possible excitation of that mode by transients or feedforward. I suggest that a notch or comb filter in the forward path is the preferred solution.

We held a code development meeting on Dec. 18 immediately following the design review. During this meeting we made a commitment to reutilize the code developed at LBNL to the extent possible, and to add new functionality on top of the existing code. Presently Craig Swanson is converting Larry's Verilog code for use on the Virtex II that will be used on the new DFE. This work is expected to be complete near the end of January. The intent is to have code ready and waiting when the prototype hardware is

completed. Mark Prokop has been working with Larry to better understand the existing code.

The filter design is under investigation. Hengjie and Larry have been discussing a notch filter concept. This is an issue for the superconducting linac only and has no impact on the testing and commissioning of the normal conducting sections.

---

2. SPECIFICATIONS

Doolittle, Champion, Ma, Regan

4. Clock: As has been recommended in the past I would come up with slow/ fast drift numbers and put bandwidths onto them. I think this was done but it was not obvious. I would also put every thing into ps. With as many frequencies as you have using degrees can be confusing. Using ps will normalize everything.

6. Closed Loop Bandwidth: This was only loosely discussed in reference with the system latency. I would like to see a Bode Plot showing the gain and bandwidth needed for field control for the various cavities (nc and sc) under a typical environment and operational parameters. This bandwidth (with contingency) should drive what system latency you must have and therefore influencing the firmware and hardware options.

Re: SNS LLRF Platform Specifications

1. VXI bus is clearly acceptable, politically expedient, and the choice that has been made. Calling it a "requirement" is out of line.

2. Please clarify "max" and "nominal". How does it relate to full scale and damage threshold? Are these levels adjustable with a screwdriver, soldering iron, or software?

See updated AFE spec (Appendix).

3. Please separate the specs for the 50 MHz and 402.5/805 MHz input channels. The needs and implementation are too distinct to share.

See updated AFE spec.

3a: Bandwidth is wrong -- should be 41 to 59 MHz

See updated AFE spec.

3e: -10 dBc output of mixer? -70 dBc output of filter?

The updated AFE spec. calls for harmonics at less than –60 dBc. This is at the output of the filter. Additional filtering on the DFE may reduce this to about –70 dBc. This is especially important for the third harmonic at 150 MHz.

4. (Output) overspecified, and I/Q terminology is obsolete. The output stage of a feedback loop is not the place for tight specifications.

See section 6.

5. (Clock) <1ms jitter -> noise on the ADC
>16ms drift -> can be corrected in software

The clock spec is revised as follows:

40 MHz ADC Clock Based on LBNL clock generation scheme.

a. Locked to 50 MHz Reference IF input.
b. Jitter < ±1 ps from 30 Hz to 500 Hz
c. Jitter < 1 ps RMS from 500 Hz to 20 MHz

8. This wish-list of memory allocation overconstrains hardware design. Can shoehorn functionality into whatever is available, above some soft minimum in the 20 kByte realm. The 6 kByte in the current LBNL design is arguably inadequate for the long term.

8l. If you want a loop delay specification, list it for pure proportional feedback. Fancier application firmware (filters) can slow this down, but that tradeoff is adjustable on-line. Separately, someone needs to investigate other filter possibilities (besides linear phase FIR bandpass) that could do the job with lower group delay.

1. **Specifications** The design specifications need to be top-down where the error budget is allocated throughout the system and the individual designs are specified according to that error budget.
    a. I recommend separating the specifications for short-term jitter and transient errors from the specifications for long-term drift.
    b. The comprehensive list of error sources should be evaluated thoroughly. At present, I believe that many error sources have not been adequately specified.
    c. All designs should be thoroughly validated according to the required specifications before being released for manufacturing or revision.

3. REFERENCE SYSTEM

Champion, Doolittle, Ma

2. Remote Mixer: While this was not presented it is one of the most important aspects of this design. I would like to see its specification like the AFE. There is still confusion as to where this lives!

We plan to use the Mini-Circuits ZFM-4H Level 17 coaxial mixer for downconversion of the reference and cavity field signals. The mixers will be mounted in a temperature regulated chassis in the LLRF control rack in the klystron gallery. Chip Piller has made measurements on the differential phase characteristics of 3/8" phase-stable Heliax over a temperature range of 70 to 100 deg F. The results support placing the mixers in the klystron gallery. This concept will be tested on the 402.5 MHz systems during DTL1 and DTL3 testing and commissioning in the spring of 2003. There are two fallback positions: 1) there is plenty of time to move the mixers to the tunnel if necessary; 2) we are pursuing the idea of time domain multiplexing the RF reference and cavity field signals and transporting them over a single path from the tunnel to the DFE.

5. Master Reference and Distribution: While this was not explicitly presented it is critical that this be included in the LLRF system. You are only as good as the master reference and some one should look at the jitter/drift as it emanates from the master oscillator out to the LO's at each mixer.

3. **RF Distribution** The RF distribution is a key part of the LLRF system and must not be ignored because it is not strictly the responsibility of LANL LLRF.

We appreciate the importance of the frequency reference and distribution on the overall system performance. The control system error budget is being revisited and will include this subsystem.

---

4. DIGITAL FRONT END (DFE)

Stettler, Doolittle, Swanson

SNS LLRF Digital Front End Specifications (Matt Stettler)

(was there a paper handout for this, other than the schematic and pin list?)

There is no document that describes the features of the DFE in detail. I will generate one before the end of the month.

- Say on paper that the FPGA program comes in via JTAG from the host.
  The XC2V1500 needs 5,166,240 configuration bits pumped over this link.
  At approximately 1 bit/usec (using the existing VXI interface), that's
  5 seconds of clock time at boot. The LBNL system, which uses bit

operations local to the CPU chip, takes approximately 1 second to load 1,040,128 configuration bits.

No problem, this is the intended method of configuration. It will be specified as the default in the DFE design description document.

- Triple check that the mechanical spec on the connectors "allows" multiple connectors on a single daughter card. (Probably OK since they have individual pins.)

I checked with Samtec, and they have never heard of such difficulties with this connector using the mechanical spacing of our board. The through hole mounting and pin/socket design makes them very robust in the face of horizontal force (typically the PC card flexes).

- Highly recommend serial number on both DFE and AFE, basically any "FRU" (field replaceable unit) with analog components.

Support has been added for an SPI bus to the AFE, DFE, and RF output boards.

- Power supply: recommend reducing output current requirement, for the first rev at least. Focus on the needs of an XC2V1500.
Can a switching regulator be found that has a synchronization input?
Plan for filtering beyond that provided within the module.

I will add support on the DFE for a pair of 3 amp linear regulators, which will be sufficient for more modest XC2V1500 designs. I will search for suitable low noise switchers - note that "medical quality" unit are usually available, I am still looking into this option. In an case, the board will retain ability to mount either the switcher or the linear regulators.

- Abandon support for the XC2V250 - it's smaller than the XC2S150 I use in the current LBNL board (3072 logic cells vs. 3456).

Done. Support for the 456 pin package has ben eliminated.

- Missing "system ADC". I see two crucial analog measurements:
VCO control voltage, FPGA core supply current.

The present design does not include a "system ADC". I am not convinced that this is a critical feature, and believe that adding it should be discussed in further detail.

- Continue revisit of DAC choice. DAC902 is available in a smaller package than what the LBNL board uses. High resolution is not needed in the output stage of a feedback loop.

The DFE supports a DAC output channel of up to 16 bits. The DAC itself was moved to the RF output board. John Power is evaluating the requirements and selecting the output DAC.

- Double check SSO (simultaneous switched outputs) restrictions on FPGA; package and drive strength dependent

I have contacted Xilinx, and been informed that this is not an issue with this device and package combination. There seems to be a bit of hand waving on this topic, but I notice that Xilinx's recommended pinout for a PCI interface on this device uses a single I/O bank. The only critical area in this design is the DAC output.

4. **Digital Front End** The design concept looks good and seems to be proceeding. Areas to be careful with include:
    a. Proper ground, power and signal isolation for low-noise design
    b. We prefer buried stripline routing surrounded by ground planes and via fences for low-noise analog signals
    c. We prefer linear regulators for ADC supplies to reduce switching noise mixing into IF signal
    d. We use differential signaling wherever possible for low-noise high-speed designs
    e. Be careful to select proper amplifiers to not degrade SINAD

a. Much effort was put into these issues on the original BPM DFE. One of the reasons we chose to use it as the basis for the new design is the performance achieved on the BPM.

c. The existing BPM design achieves 2 bits RMS noise on the 14 bit ADCs; this performance is consistent with the DFE requirements.

d. In our design, we will be using differential signaling on critical clock distribution nets.

e. It is important that the amplifiers in the analog section be chosen with care. There are no amplifiers on the DFE on the ADC inputs. The output to the RF output board is digital in the current design.

---

5. ANALOG FRONT END (AFE)

Power, Doolittle, Ma

1. Analog Front End: Although I have not seen the board specification (what you send to Bergoz) what was presented Tuesday was some what confusing. Ranges need to be added for SNR (over what dynamic range etc.). Isolations and filter specifications are awkward. I would add an IP3 specification for the various

channels (mixers). Some of the specifications may be redundant/conflicting especially as they concern distortion and isolation. An acceptance test plan as well as test procedures for some of the more subtle measurements needs to be made. Are the BPM switches needed? If you have a loop back feature great. It has been my experience that switches lower board isolation. John Power should continue his System View models looking at the various channels.

SNS LLRF Analog Front End Specifications (John Power)

7. (drift) Distinguish between absolute and differential phase drift
Mention amplitude sensitivity of phase measurement (AM/PM conversion)

2. **Analog Front End** In general, the specs look fine. I suggest separating the non-critical signal specs (FWD, REF) from the critical signal specs (IF, CAV). For the critical signals, the following issued must be considered:
    a. Tight specification of SINAD over signal band of interest – do not over-specify for out-of-band signals
    b. Mixer IP3 and sensitivity to AM-to-PM
    c. IF filter return loss at IF and image frequencies
    d. IF input return loss and interaction with IF filter output return loss (or whatever is on far end of cable)

Please see the updated AFE specification, which is included as an appendix to this document.

---

6. ANALOG OUTPUT

Power, Doolittle, Ma

3. Analog Output: I suppose that this is an up converter and a switch now??

The RF Output (RFO) board is responsible for taking the 50 MHz baseband signals from the Digital Front End (DFE) and upconverting them to the RF frequency. In addition it will be the home for the phase locked loop clock generation circuitry, which is the basis for all of the coherent sampling. Essentially the RF Output board consists of a 14-bit DAC followed by the 50 MHz Output Channel circuitry used in the LBNL design. This will be the input to an Analog Devices active mixer AD8343. The mixer output will be bandpass filtered and routed to the front panel. The LO will come in through a splitter, one leg of which will go right back out to the AFE as an input. The other leg will be attenuated and fed into the LO channel of the AD8343 mixer

At the time the DFE design was deemed ready for ECAD, John had not decided which DAC to use, so rather than hold up the ECAD work, we decided to put the DAC on the RFO board. The Diagnostics BPM DFE required revisions in order to minimize the impact of switching noise between the ADCs and the FPGA. John and Matt were

concerned about similar interference problems at the output end of the FPGA. So they decided to put the DAC on the RFO where the analog ground will have small series resistors to stop any ground currents from flowing through the signal lines, similar to what was done on the ADC channels for the BPM DFE.

The analog output specification is revised as follows:

Analog Output Channel Specification

a. 30 dB dynamic range (-20 to +10 dBm)
b. Resolution: 0.2 % amplitude, 0.15 deg phase
c. Linearity < 5 %
d. Spurious and Harmonics < -60 dBc

APPENDIX

SNS LLRF Analog Front End Specifications:
1/29/03

All performance specifications apply to nominal operating conditions unless otherwise noted.

1. Physical Architecture

1. Identical to footprint for Beam Position Monitor presently being provided by Bergoz, (although SMA connectors can be moved slightly if necessary). Overall board size and edge connector location should not change. See below.

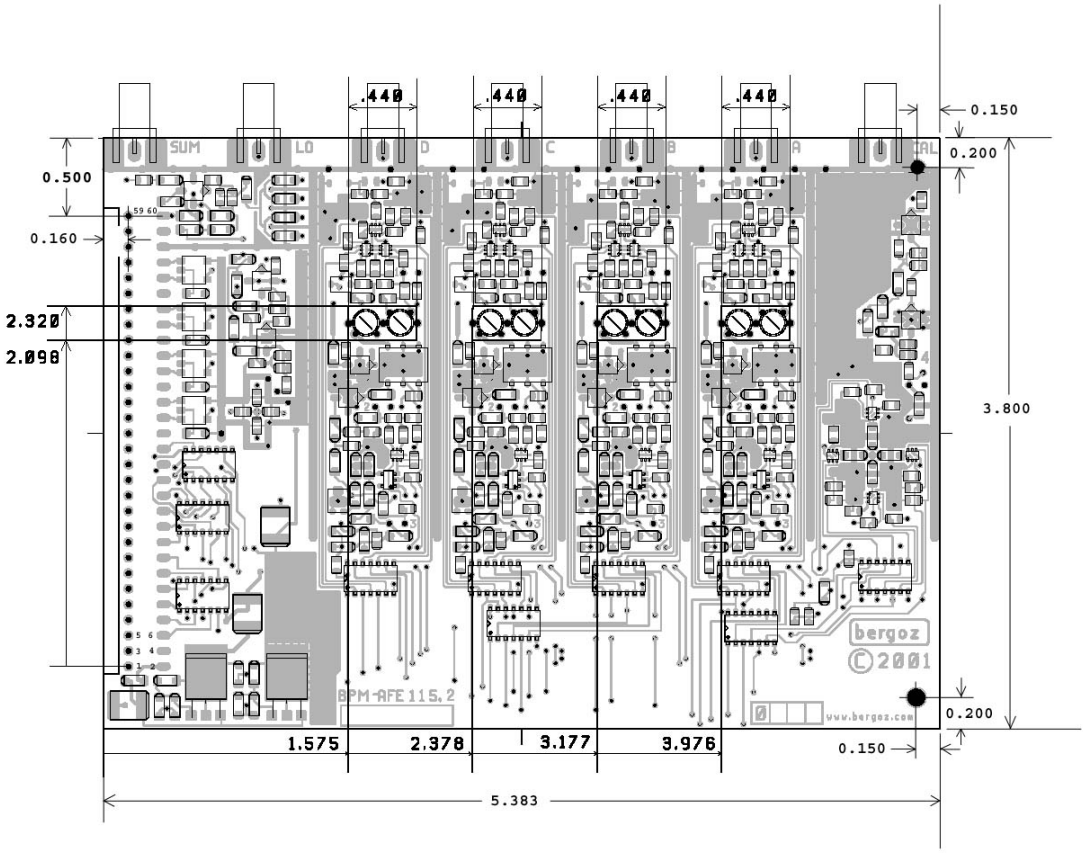

2. Input channel connectors: SMA edge connectors as per BPM AFE
3. Output connector: as per BPM AFE. See below.

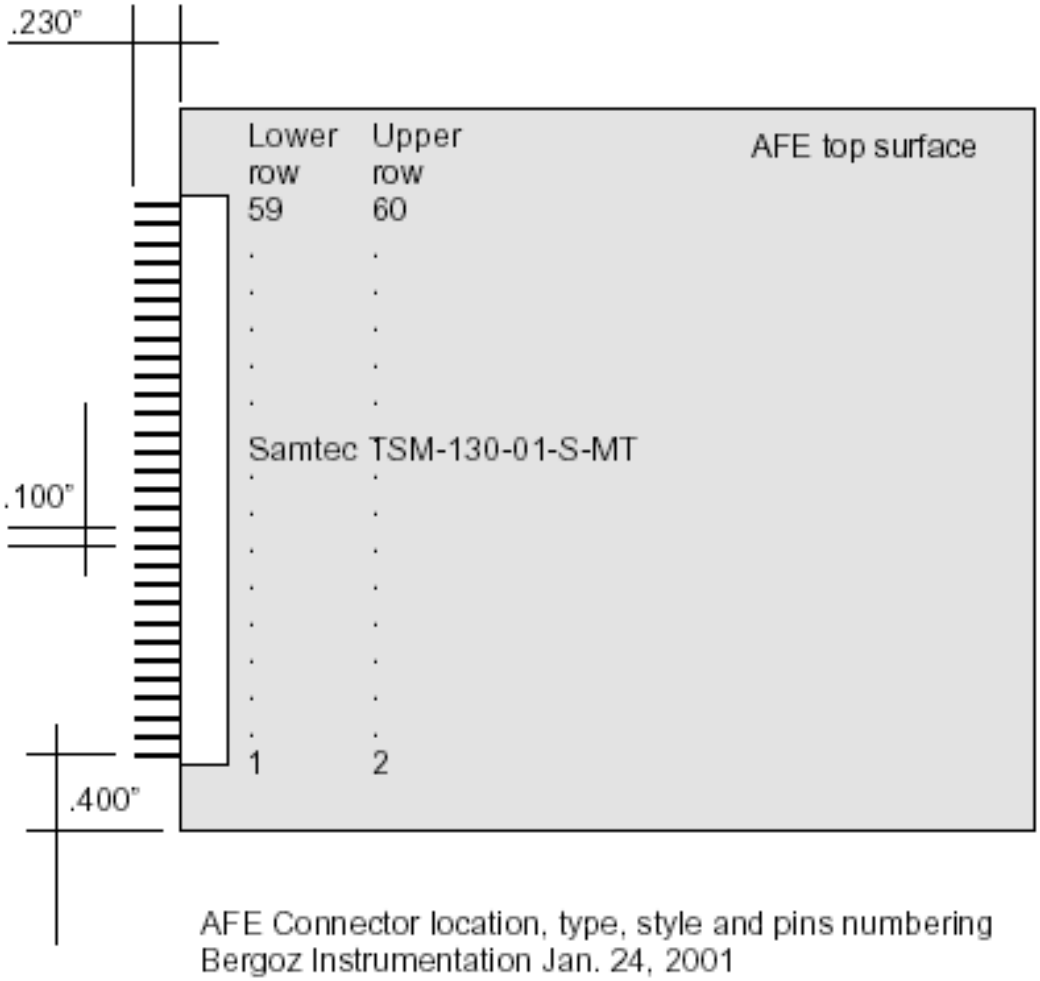

AFE Connector location, type, style and pins numbering
Bergoz Instrumentation Jan. 24, 2001

Expected pinout for the 60 pin connector, per the diagram above, is:

1 GND	2 +12V
3 +5V	4 –12V
5 SCLK	6 A.PICKUP-SW
7 SERcs2	8 A.RCVR-SW
9 +3.3V	10 A.CAL-SW
11 GND	12 B.PICKUP-SW
13 GND	14 B.RCVR-SW
15 GND	16 B.CAL-SW
17 GND	18 C.PICKUP-SW
19 GND	20 C.RCVR-SW
21 GND	22 C.CAL-SW
23 GND	24 D.PICKUP-SW
25 GND	26 D.RCVR-SW
27 GND	28 D.CAL-SW
29 GND	30 SPARE (FPGA)
31 GND	32 GND
33 GND	34 A.OUT-

35 GND	36 A.OUT+
37 GND	38 GND
39 GND	40 B.OUT-
41 GND	42 B.OUT+
43 GND	44 GND
45 GND	46 C.OUT-
47 GND	48 C.OUT+
49 GND	50 GND
51 GND	52 D.OUT-
53 GND	54 D.OUT+
55 GND	56 GND
57 GND	58 SDI1
59 GND	60 SDI0

2. Analog Input Connections

1. Forward RF (402.5/805 MHz) +11 dBm max, +10 dBm nominal.
2. Reflected RF (402.5/805 MHz) +11 dBm max, +10 dBm nominal
3. Cavity IF (50 MHz) +3 dBm max, +2 dBm nominal
4. IF Reference (50 MHz) +3 dBm max, +2 dBm nominal
5. LO (352.5/755 MHz) +6 dBm max, +3 dBm nominal

Items 1,2,3: max implies headroom before ADC saturation (e.g., ADC saturates above an +11dBm Forward RF signal level).

3. IF Channel Specifications (Cavity and Reference)

1. Damage threshold input level = +20 dBm.
2. SNR 60 dB per channel minimum, integrated over a 41 MHz to 59 MHz bandwidth.
3. Isolation between IF Field and IF Reference channels at 50 MHz: 70 dB minimum
4. Signal distortion, each channel: 60 dB relative to 50 MHz, for the first four 50 MHz harmonics.
5. |S11| shall be < -25 dB (30-70 MHz) on the 50 MHz IF inputs.
6. Field Sample/Reference IF input to IF output amplitude linearity: +/- 0.5% deviation from linear.
7. Field Sample/Reference IF input to IF output phase linearity: +/- 0.5 degree deviation from linear.
8. Each differential IF output, at nominal input level, shall be $\pm 1.1\ V_{p\text{-}p}$ into 50Ω.

4. LO Channel Specification

1. Damage threshold input level = +17 dBm.
2. |S11| shall be < -20 dB over a ±5 MHz passband centered at 50 MHz.

5. RF Channel Specification (downconverting channels – Forward and Reflected RF)

1. Damage threshold input level = +17 dBm.
2. SNR 60 dB per channel minimum, integrated over an RF ± 9 MHz bandwidth.
3. Isolation between forward and reflected channels at 402.5/805 MHz: 45 dB
4. Both RF channels shall downconvert the RF input to a 50 MHz output, utilizing the LO input.

5. The Forward RF and Reflected RF channels each require a ±20 MHz bandwidth band pass filter centered at 402.5/805 MHz on the inputs before the mixers.
6. All mixer harmonics need to be down by at least 40 dB relative to the fundamental 50 MHz output.
7. Signal distortion, each channel: 60 dBc relative to 50 MHz, for the first four 50 MHz harmonics.
8. |S11| shall be < -20 dB over a ±5 MHz passband centered at the 402.5 or 805 MHz RF frequency for both the forward and reflected channels, -20 dB (LO freq ±5 MHz) on the LO.
9. Forward/Reflected RF input to IF output amplitude linearity: +/- 0.5% deviation from linear.
10. Forward/Reflected RF input to IF output phase linearity: +/- 0.5 degree deviation from linear.
11. Each differential IF output, at nominal input level, shall be ±1.1 V_{p-p} into 50Ω.

6. Power

Unipolar +12 V

7. Drift

Expected operating temperature range: 67 - 100 °F
Phase drift: ±1.04 ps max (±0.3 deg at 805 MHz)